

Pilot Interactions in an Over-constrained Conflict Scenario as Studied in a Piloted Simulation of Autonomous Aircraft Operations

David J. Wing *

Dr. Karthik Krishnamurthy †

Richard Barhydt *

Dr. Bryan Barmore *

* NASA Langley Research Center, Hampton VA USA

† Titan Corporation, Hampton VA USA

Abstract

Feasibility and safety of autonomous aircraft operations were studied in a multi-piloted simulation of over-constrained traffic conflicts to determine the need for, and utility of, priority flight rules to maintain safety in this extraordinary and potentially hazardous situation. An over-constrained traffic conflict is one in which the separation assurance objective is incompatible with other objectives. In addition, a proposed scheme for implementing priority flight rules by staggering the alerting time between the two aircraft in conflict was tested for effectiveness. The simulation was conducted in the Air Traffic Operations Laboratory at the NASA Langley Research Center. This research activity is a continuation of the Distributed Air-Ground Traffic Management feasibility analysis reported in the 4th USA/Europe Air Traffic Management R&D Seminar in December 2001 (paper #48).

The over-constrained conflict scenario consisted of two piloted aircraft that were assigned an identical en-route waypoint arrival time and altitude crossing restriction. The pilots safely resolved the conflict without the need for a priority flight rule system. Occurrences of unnecessary maneuvering near the common waypoint were traced to false conflict alerts, generated as the result of including waypoint constraint information in the broadcast data link message issued from each aircraft. This result suggests that, in the conservative interests of safety, broadcast intent information should be based on the commanded trajectory and not simply on the Flight Management System flight plan, to which the aircraft may not actually adhere. The use of priority flight rules had no effect on the percentage of the aircraft population meeting all assigned constraints. However, the

priority system made completely predictable which aircraft in a given pair would meet the constraints and which aircraft would make the first maneuver to yield right-of-way. Therefore, the proposed scheme for implementing priority flight rules through staggering the alerting time between the two aircraft was completely effective. The data and observations from this experiment, together with results from the previously reported study, strongly support the feasibility of autonomous aircraft operations.

Introduction

The NASA Advanced Air Transportation Technologies project is conducting exploratory research and development on a far-term concept of operations for Air Traffic Management (ATM) defined by a redistribution of ATM responsibilities between air traffic service providers and aircraft flight crews. The operational concept is called *Distributed Air/Ground Traffic Management* (DAG-TM)^[1], and many of its elements proceed along the conceptual path offered by the original RTCA Free Flight concept^[2] wherein flight crews select their path and speed in real time while conforming to restrictions established for safety and flow management. One of the DAG-TM concept elements^[3] describes operations in the en-route and terminal-transition domains and establishes a clear delineation of responsibilities between the groundside and airborne participants within these domains, albeit a significant shift from responsibilities in current-day operations. The principal shift proposed in the concept element is that properly trained flight crews of properly equipped aircraft assume full responsibility for separation from similarly equipped traffic throughout the en-route and terminal-transition domains. Aircraft not in this category continue to receive separation services from the

ground. The primary anticipated benefit of creating this new category of aircraft operations is the ability of the National Airspace System (NAS) capacity to dynamically adapt to significant variations in demand, thereby accommodating a substantial increase in traffic volume over that manageable by a ground-based system. This mechanism towards increased capacity is referred to as “scalability” and would likely result from minimizing the interactions between this new category of “autonomous aircraft” and the ground-based Air Traffic Service (ATS) provider. With the interactions minimized, it is hypothesized that the volume of self-responsible traffic could grow nearly independently from ground-based operations and infrastructure. Not surprisingly, minimizing the interaction is the principal challenge of the concept development. The concept element includes many features designed specifically for this purpose. An example is the full responsibility of autonomous aircraft both for assuring separation with similar aircraft and for conforming to flow constraints assigned by the ATS provider, with intervention by the ATS provider neither required nor expected in these tasks.

Proceeding on the assumption that the interactions can indeed be minimized, it is acceptable to study aspects of autonomous aircraft operations in isolation from the operations of the ATS provider, provided the areas where interactions do occur are avoided or carefully handled. Given that ground-based separation services for autonomous aircraft are removed from the equation, the principal interaction under nominal conditions involves the second fundamental role of the ATS provider, that is, Traffic Flow Management (TFM). With traffic volume envisioned to grow significantly in the coming years, congestion in terminal areas will certainly also grow and will likely expand to cover most hours of each day. TFM will therefore become a continual operation for the ATS provider, rather than an intermittent task. The DAG-TM concept states that arrival metering will be the principal TFM tool, and if needed, en-route metering as well. Armed with predictive information on arrivals and NAS status, the ATS provider establishes the arrival sequence and meters the flow into the terminal area by issuing required-time-of-arrival (RTA) clearances and crossing restrictions at inbound metering fixes. Once these clearances and restrictions are received and accepted by a given flight crew, the interaction between the ATS provider and this autonomous aircraft are thereafter minimized until the aircraft crosses the fix. At this time, the aircraft commences terminal arrival operations, which the DAG-TM concept treats as a separate concept element with different assumptions on roles and responsibilities^[4].

In addressing concept feasibility, the basic scenario of concern for the autonomous aircraft is the ability of the flight

crew, after having received and accepted the arrival clearance, to plan and execute an efficient conflict-free trajectory to the metering fix, arriving on time and altitude within established tolerances. Barring for the time being other interactions with the ATS provider, of which conflicts with ground-managed aircraft are an example, this airborne scenario can be investigated in simulation without the inclusion of the ATS provider role, thereby greatly simplifying the simulation environment and experimental logistics. Several research questions might be asked regarding the feasibility limits of this scenario. What metering interval between arriving autonomous aircraft is too small for aircraft to self-separate while converging to the metering fix? How close to the metering fix can autonomous aircraft receive and adjust to a revised crossing time, while still remaining separated from traffic (i.e., the issue of the RTA freeze horizon)? These and other issues define the limits of acceptability on the constraints imposed by the ATS provider. To minimize the air/ground interactions, these limits need to be well understood. For example, schedule the aircraft too closely or change the constraints too late, and the distributed system of autonomous aircraft may be unable to reliably meet the constraints while maintaining separation. Yet provided that these limits are well understood, the ATS provider can ensure that the constraints imposed upon autonomous aircraft are achievable.

Previous research on this DAG-TM concept element found no feasibility impediments to autonomous aircraft meeting *achievable* operational constraints under nominal conflict situations^{[5],[6]}. In reporting on DAG-TM concept feasibility, it may be sufficient to rest on the statement that any operational constraints imposed upon autonomous aircraft must be determined to be achievable before they are assigned. However, concept robustness requires the investigation to be extended beyond near-nominal conditions. Taken to an extreme condition, how will the DAG-TM concept behave in a situation where it is literally impossible for all constraints to be simultaneously met? Will safety be compromised if pilots are unaware of the over-constrained nature of the situation? Are there protective measures that can be incorporated into the airborne side of the operational concept such that additional air/ground interaction is not necessary to manage this anomalous event? If the concept can be shown to be robust under such extraordinary situations, then the feasibility argument is strengthened, because the assumption of assigning only achievable constraints will have been shown to not be of such high criticality.

The experiment reported herein addressed this issue in a human-in-the-loop simulation of autonomous aircraft operations. The exploratory study attempted to determine the

nature of pilot interactions in an over-constrained conflict situation and the need for and effectiveness of priority flight rules in preserving traffic separation. The experiment was performed in the NASA Langley Research Center Air Traffic Operations Laboratory, a distributed desktop simulation of aviation operations in which pilots of multiple simulated aircraft can interact in preplanned or dynamically developing scenarios using prototype decision support tools and procedures under development for DAG-TM operations. This research activity is a continuation of the DAG-TM feasibility analysis reported in the 4th USA/Europe Air Traffic Management R&D Seminar in December 2001^[5].

Over-constrained Conflict Scenario

The experimental scenario chosen to investigate over-constrained conflict situations is depicted in Figure 1. Two autonomous aircraft, each controlled in the experiment by a separate subject pilot, are established on flight plans that intersect at a common waypoint. To meet hypothetical TFM needs (details not included or required in the simulation), the flight plan of each aircraft at this waypoint includes an RTA and a “cross at” altitude restriction, both of which were supposedly assigned by the ATS provider. These constraints are pre-entered in the Flight Management System (FMS) flight plan at the outset of the simulation, which begins approximately 200 miles (25 minutes) before the common waypoint.

From an operational standpoint, this common waypoint could be considered an en-route or terminal-arrival metering fix. En-route metering has been hypothesized as a TFM technique for regions of highly constrained special-use airspace (SUA), weather, or traffic congestion, as well as for

absorbing arrival delay in the en-route domain^[1]. Terminal-arrival metering is in limited use today as an arrival flow management technique^[7]. In the latter case, the crossing restriction is typically established well below cruise altitude (i.e., below 14000 feet). The early prototype version of the airborne conflict management tools used in this simulation lacked sufficient vertical functionality to allow this scenario to be studied in descent. Nevertheless, the objectives of the experiment could be met with the entire scenario flown in level cruise flight.

To create the over-constrained conflict, identical time and crossing altitude assignments were given to both aircraft, creating an impossibility for both aircraft to exactly meet these constraints and simultaneously maintain separation. For such an event to occur, of course, a significant failure must have occurred at some level in the NAS, and probably multiple failures. For example, the ground system could have erroneously assigned the same RTA, a data link or voice transmission error could have occurred, or one of the pilots could have entered the constraint incorrectly into the FMS. Additionally, any ground systems in place to monitor predicted arrival times at the fix and detect problems would have had to malfunction. Despite the probable rarity of these failures, the interest remains on how the airborne side reacts and responds to such a potentially hazardous situation. In running the experiment, an effort was made to ensure neither pilot knew that (a) their conflict was with another test subject rather than the typical computer-driven scripted aircraft and (b) both aircraft shared a common set of crossing constraints.

Prior to the common waypoint, the scenario for each aircraft included passage through a corridor defined by two SUA's, through which many additional scripted aircraft also flew. This aspect of the scenario was included primarily to support other experimental objectives outside the scope of this paper, but also to provide a realistic and distracting environment for the pilot in advance of the planned conflict.

Airborne Coordination

The DAG-TM en-route concept element simulated in this investigation includes neither addressed data link nor voice communication between aircraft. All information flow between aircraft therefore occurs only through broadcast data link transmissions. The quantity and type of information transferred will almost certainly be constrained by technology and bandwidth limitations^[8], and the message content will therefore be likely to only include basic information critical to general situation awareness and conflict detection such as aircraft identification, equipment, current state, and limited intent. If it can be shown that non-explicit coordination between aircraft through this limited broadcast information set allows safe resolution of conflicts,




Figure 1. Traffic scenario in which flow management constraints conflict with separation constraints.

it would then be unnecessary to create additional technical complexity to enable aircraft to coordinate explicitly. The feasibility analysis challenge, therefore, is to determine whether non-explicit (implicit) coordination is satisfactory. If it were to be determined that an extreme situation such as an over-constrained conflict does not require explicit coordination, then nominal situations are unlikely to require it either.

Some forms of implicit coordination do occur, however, with broadcast data link, as envisioned in this concept element. The message set includes current aircraft state and in most cases some limited intent information. Once an aircraft has detected a conflict, a change in either the state or intent information in its broadcast message most likely indicates an awareness of the conflict and the intention to resolve it. If the change includes new intent, then the method by which the resolution will occur is also made clear. If the conflict is near-term, then tactical resolution advisories are displayed to each crew. The compatibility of these maneuver advisories is being ensured by using a common set of advisory algorithms. Only implicit coordination methods such as these were included in the simulation, and this choice was deliberate to meet the feasibility challenge stated earlier.

Priority Flight Rules

Apart from coordination through data exchange, it is also possible to coordinate a response to a conflict through a set of common procedures. A rule base could be established that defines the relative priority between aircraft in conflict, and as long as both aircraft use the same rule base, implicit coordination has effectively occurred without any data exchange. The aircraft with the lower priority would be expected to assume the burden of resolving the conflict. This is similar to the use of Visual Flight Rules (VFR) “right of way” rules.

So is implicit coordination sufficient to preserve safety in an over-constrained conflict situation? If so, is the assignment of relative priority between the aircraft a necessary component of implicit coordination? The experiment attempts to gain insight to these questions by studying how the chosen scenario plays out with and without priority flight rules in effect.

Defining which aircraft has priority in a given conflict situation can easily become a complex process. The operational concept states that autonomous aircraft have priority over managed aircraft. Yet further study of this issue indicates some exceptions to this rule that may be warranted under certain conditions involving differing time horizons and auto-flight modes. In conflicts between two autonomous aircraft, the conflict geometry can be used to determine priority, such as the VFR stipulation that the right-hand

aircraft has right-of-way. Yet this too is complicated by factors such as distinguishing head-on and overtaking tracks from crossing tracks. Other factors have been proposed to drive priority that further complicate matters. Examples include whether aircraft are in cruise, climbing, or descending flight, whether aircraft are broadcasting full or reduced intent, and whether aircraft are on or off schedule to a near-term RTA.

If all of these considerations must be included in the priority calculation, it would be unreasonable to expect pilots (or controllers) to remember and accurately apply these distinctions and thereby ensure implicit coordination. A method was therefore developed to implement priority flight rules such that the actual rule set is transparent to the pilot and any level of complexity can therefore be incorporated. The implementation method takes advantage of a multi-stage alerting scheme defined by the Airborne Conflict Management (ACM) working group of RTCA Special Committee 186^[9] and implemented with enhancements for this experiment. Four alert levels were used regarding traffic information as displayed to the flight crew. A Level 0 alert, the lowest level, is essentially a traffic point-out to assist in pilot awareness (i.e., no conflict may exist) and the flight crew is not required to take action*. The Level 1 alert equates to the ACM definition of the “low level alert” where a conflict exists but is sufficiently distant that responsive action by the flight crew is not required but may be encouraged by company policy. The Level 2 alert is a conflict alert for which the flight crew is legally required to take timely resolution action. The Level 3 alert addresses a possible collision threat and requires immediate action. These alerts typically appear in succession as a function of time as the conflict draws nearer. In the simulation, the look-ahead time horizon (maximum detection range) for intent-based conflicts was chosen to be 10 minutes, and for state-based conflicts, 5 minutes†. For the type of intent-based conflict described for this experiment, the conflict aircraft would typically first appear as a Level 1 alert at 10 minutes prior to predicted loss of separation (LoS) and upgrade to a Level 2 alert at 5 minutes, at which point the flight crew would be required to take action. Prior to the Level 1 alert, the conflict aircraft would not be displayed at all, unless it triggered a Level 0 alert for some other reason.

If priority rules are not enacted, then both aircraft would receive each alert level at approximately the same time.

* The Level 0 alert was an addition to the ACM working group alerting design. Aircraft that do not meet the criteria for a Level 0 alert are not displayed at all.

† Intent-based conflicts are detected using best-available trajectory information from both aircraft, while state-based conflicts are detected using only state data, with no regard for availability of intent information.

There is a possibility that one or both aircraft would attempt to resolve the conflict between 5 and 10 minutes prior to LoS, and a high probability (assuming the pilots adhere to their responsibility) that both will take action at 5 minutes. Of course, any resolution action will initially result in abandonment of the assigned operational constraints at the waypoint. If both aircraft simultaneously maneuver to resolve the conflict, will one or both pilots observe this and subsequently attempt to reestablish their aircraft toward meeting the assigned constraints? If so, then a dangerous dynamic could be established as each aircraft proceeds with the assumption that the other aircraft will give way.

To implement priority flight rules, an approach has been developed whereby the alert levels presented to the pilots are staggered relative to each other according to the calculated priority. The resulting succession of alerts is depicted in Figure 2. When a conflict is detected by the flight deck decision support system, the alerting algorithm considers all factors relevant to aircraft priority, as described earlier, and assigns either a higher or lower priority relative to the conflict aircraft. The appropriate alert level is then displayed. The resulting effect is that, up until the 2-minute point, the pilot of the lower-priority aircraft is viewing a higher alert level and is therefore receiving a greater encouragement to resolve the conflict. It is hypothesized that this bias spurs the lower-priority aircraft to maneuver first. Furthermore, as a result of the implicit coordination through the change in broadcast intent, the higher-priority aircraft is induced to take no action, as the conflict alert will have disappeared. Then with no change in the higher-priority aircraft's broadcast intent or state, the lower-priority aircraft would presumably be less likely to renew the effort to achieve the assigned constraints.

This alerting scheme may also assist in simplifying resolution of multi-conflict situations (where one aircraft is simultaneously in conflict with two or more aircraft). Referring to Figure 2, if a third aircraft were in conflict with the two shown, and it had a relative priority in between the other two, the Level 2 alert could be displayed to this pilot at a time between 5 and 2 minutes to LoS, thereby influencing the order of maneuvering.

Experiment Approach

A human-in-the-loop experiment was devised to investigate the need for and effectiveness of using priority flight rules to maintain safety in an over-constrained conflict situation involving two autonomous aircraft. The experiment involved the participation of 16 active or recently active airline pilots flying desktop flight simulators in an interactive future airspace environment. Only the autonomous aircraft operations aspect of the DAG-TM en-route concept element

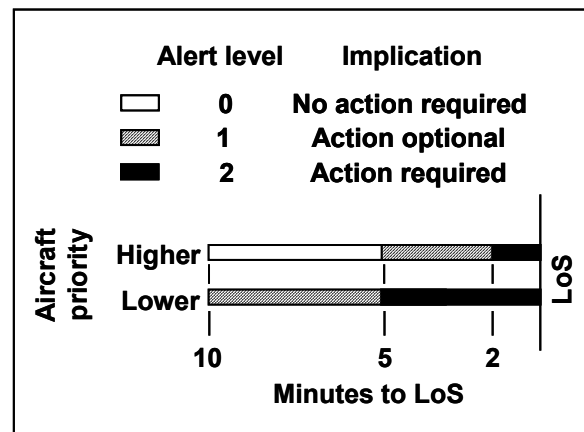


Figure 2. Implementation of priority flight rules through staggered alert levels.

was represented in the simulation; ATS provider control of managed aircraft was not simulated, nor was it needed to achieve the experimental objectives.

Four pilots flew simultaneously in the simulation, and efforts were taken to disguise the fact that pairs of them would occasionally interact through traffic conflicts. The pilots received comprehensive training on the procedures, alerting levels, and conflict management tools through printed material, classroom briefings, and hands-on practice with one-on-one instruction before they flew the planned test scenarios. Additionally, the first test scenario after training was in fact a buffer scenario to verify the required pilot proficiency level had been reached, unbeknownst to the pilots. Including this buffer scenario, each pilot flew 10 scenarios, each based on the same basic scenario design depicted in Figure 1. Six of these scenarios contained conflict situations designed to investigate separate experimental objectives outside the scope of this paper, and they did not contain the over-constrained conflict where pairs of aircraft were given the same crossing assignment at the same waypoint. The remaining four scenarios did contain the over-constrained conflict, including the buffer scenario that was not intended for inclusion in data analysis other than to verify pilot proficiency.

The three scenarios relevant to this paper included the following conditions and pilot-pair counts[‡]:

1. No priority flight rules in effect. (N=10)
- 2a. Priority flight rules in effect, odd-numbered pilot has priority. (N=8)
- 2b. Priority flight rules in effect, even-numbered pilot has priority. (N=4)

The test matrix can also be considered to have just two conditions: priority rules not in effect (N=10), and priority

[‡] Pilot-pair counts are unbalanced across conditions due to simulation faults and lost data during the experiment runs.

rules in effect (N=12). All three scenarios paired the same two pilots. Overall, the nine test scenarios (excluding the buffer scenario which was always the first scenario) were counterbalanced between experimental objectives, of which this paper addresses only one.

The pilots were each instructed to maintain traffic separation as a top priority, and as a second priority, to achieve the assigned waypoint constraints. Acceptable waypoint crossing tolerances were to cross within 2.5 nautical miles (nm), within 500 feet (ft) of the assigned altitude, and within 30 seconds of the RTA. A map display depicting the waypoint and the arrival time error, as well as an autoflight system capable of meeting the RTA, were provided for assistance. If the pilot determined that one or more of these waypoint constraints could not be achieved for any reason, the instructions were to “notify” the ATS provider of this fact at the earliest possible time. Three buttons provided in the flight simulator to record this notification were labeled “unable fix,” “unable alt,” and “unable time” (in the absence of an actual ATS provider and air/ground communication in the simulation). This notification would presumably constitute the first step in the process of requesting a new metering assignment. In addition to these tasks, the pilots were prompted every one to two minutes with a secondary distracting task of answering aviation and trivia questions.

The minimum traffic separation criteria were to avoid flight within another aircraft’s cylindrical protected zone with dimensions of 5 nm radius and ± 1000 vertical feet. Conflict prevention, detection, and resolution tools were provided to assist the pilot in accomplishing this task. One such tool set, the Autonomous Operations Planner (AOP), is under development for research purposes at the NASA Langley Research Center^[10]. The current experiment used an early prototype version of AOP^[11] (containing incremental enhancements to the version described in the reference). It provided capabilities for conflict detection and alerting that consider both state and intent information and conflict resolution advisories supporting both tactical and strategic maneuvers[§]. Intent information was broadcast through a simulated Automatic Dependent Surveillance Broadcast (ADS-B) network. This information was used for intent-based conflict detection and included flight plan Trajectory Change Points (TCPs) and any flight plan constraints such as the RTA or altitude restrictions at waypoints present in the FMS flight plan.

Recorded data included trajectories of all aircraft, all pilot actions in controlling the aircraft through the FMS and Flight

Control Panel (FCP), pilot manipulations of traffic data, alert levels and times, actuation of the “unable” buttons, and subjective questionnaire responses after the experiment. The results discussed below were based on analysis of the recorded data supplemented by researcher observations in the course of the experiment. For the most part, there were insufficient samples to perform formal statistical analyses, however qualitative analysis has yielded several important findings.

Discussion of Results

General discussion of safety

A fundamental indication of whether safety was compromised in this scenario could be considered the occurrence of LoS events. Of the 22 data runs focused on over-constrained traffic conflicts, which were designed for LoS, two runs actually contained LoS events, and as one would expect, both LoS events occurred close to the waypoint where the aircraft shared common RTA and altitude constraints. Both events occurred in data runs where priority flight rules were in effect. However, the priority flight rules did not appear to play a role in either situation.

In the first case, both aircraft had deviated from the assigned constraints, one aircraft (A) by descending 1000 ft below the required crossing altitude and the other aircraft (B) by laterally bypassing the waypoint by ~ 6 nm and passing behind the first aircraft at approximately a right angle. Aircraft A had actually descended to resolve a separate conflict. Aircraft B, the lower priority aircraft, had maneuvered to resolve the conflict with aircraft A. Upon observing aircraft B pass behind him on the Navigation Display (ND), aircraft A initiated a climb, attempting to meet the altitude constraint. However the climb was premature, since the required 5-nm separation had not yet been reached. The climb resulted in a brief LoS, albeit the collision threat was nil^{**}. While still at the lower altitude, aircraft A had declared “unable altitude constraint” by clicking the appropriate data link button, and therefore was no longer required to meet the crossing restriction. Had the ATS provider been included in the simulation and acknowledged the declaration, a new altitude assignment would have been issued, and the proximate passing of the aircraft would have been averted. So, despite the LoS event, safety was not truly compromised in this scenario.

In the second LoS case, the lower priority aircraft (C) had climbed 1000 ft above the assigned altitude prior to reaching the waypoint, thereby yielding to the higher priority aircraft (D). As the aircraft approached the waypoint, the autoflight system of aircraft C initiated an unexpected descent to meet

[§] Only horizontal strategic resolutions were available from the AOP in this simulation. Both horizontal and vertical tactical resolutions were available.

^{**} Case 1: Minimum separation was 4.67 nm and 979 ft.

the altitude crossing restriction still contained in the FMS flight plan, resulting in immediate LoS with aircraft D. The alert pilot, though surprised by the descent (as indicated in a post-scenario questionnaire), arrested the descent after 150 ft and reestablished the required separation by climbing^{††}. In this case, the sudden descent was the result of a simulation error, since the autoflight system violated the limiting altitude set by the pilot on the Flight Control Panel. Neither the pilot, nor the conflict management decision aid, was aware that the autoflight system would command this descent. This scenario underscores the importance of linking the autoflight system (including the FMS and FCP settings) to the conflict management decision aid. The conflict detection system should continually check the trajectory the autoflight system is configured to fly (the “commanded” trajectory) and provide appropriate alerting to the flight crew. Such capability is currently being implemented in an updated version of AOP^[10].

Following each data run, the pilots were given a questionnaire that included the following question: “Considering the complete start-to-end scenario (including the conflicts and your resolution actions), what was the level of safety?” The ordinal rating scale ranged from “not at all safe” (1) to “neutral” (4) to “completely safe” (7). Responses are shown in Figure 3. Out of 44 responses (22 data runs), 77% rated the level of safety higher than neutral, with half of these finding the scenario completely safe.

Although the pilots were not directly asked the reason for their rating of safety, general comments provided in free response segments of the questionnaires, by the 6 pilots who rated safety less than neutral, provide some insight to their experiences. One pilot reported frustration at having to ignore a false conflict alert, an unfortunate artifact of the simulation configuration of broadcasting waypoint *constraint* information rather than the true commanded trajectory. Two pilots reported conflict alerts but no resolution guidance, a situation that was possible with the current prototype tools under certain unique combinations of autoflight mode and conflict geometry. One pilot reported that unfamiliarity with the simulator’s autoflight system resulted in unexpected aircraft behavior (the second LoS scenario described earlier). Another pilot thought he had misunderstood the separation requirements and had therefore himself contributed to reduction in safety. The 6th pilot expressed frustration with finding a conflict at the RTA waypoint; the objective of the experiment, of course, was to determine whether such an extraordinary and potentially hazardous situation could be safely resolved. With the two exceptions described earlier, the over-constrained conflict situations were safely resolved




Figure 3. Frequency of pilot ratings of safety level across all over-constrained conflict scenarios. Ordinal scale ranged from (1) “not at all safe” to (7) “completely safe.”

(separation was preserved), and in general, the pilots agreed the operation was acceptably safe.

Who Met the Waypoint Constraints

The conflict designed for this experiment was considered over-constrained because it was not possible for both pilots to meet their assigned RTA and altitude constraints within the established tolerances while also maintaining traffic separation. Therefore one or both pilots would need to abandon a constraint, or otherwise risk a separation violation. The constraint conformance of the subject aircraft across all data runs is presented in Figure 4. Three comparisons are presented in this figure: (1) all subject-aircraft with and without priority flight rules in effect; (2) subject-aircraft approaching from the right-hand (RH) and left-hand (LH) directions when priority flight rules were not in effect; and (3) subject-aircraft with and without priority when flight rules were in effect.

Approximately one-third of the aircraft met all of the waypoint constraints (time, position, and altitude) within tolerance regardless of whether the priority flight rule system was used (1). When priority flight rules were not used (2), the RH aircraft in a conflict pair met their constraints about 2.5 times as frequently as the LH aircraft. This trend may be the result of pilots instinctively or subconsciously applying previously learned standards for priority regarding approaching vehicles. In VFR, as well as on waterways and stop sign intersections, the RH vehicle is considered to have right-of-way. When priority flight rules were in effect (3), aircraft that met all constraints were always the higher priority aircraft, indicating that using a priority system increases predictability regarding which aircraft in a pair will be likely to prevail.

^{††} Case 2: Minimum separation was 1.92 nm and 998 ft.




Figure 4. Ability of pilots to adhere to the assigned waypoint constraints.

With priority flight rules in effect (3), in only two-thirds of the cases did one aircraft of the pair (the higher-priority aircraft) meet all constraints. The expectation was that this ratio would be higher, since abandonment of any constraints by lower-priority (LP) aircraft should have permitted all higher-priority (HP) aircraft to meet all constraints. An investigation of the conflict alerting revealed that, in three of the four cases where the HP aircraft did not meet all constraints, the HP aircraft had received a conflict alert *after* the LP aircraft had already implemented strategic resolutions^{**} delaying their arrival at the RTA waypoint enough that the HP aircraft could arrive on time. The lateral path-stretch that created the arrival delay should have eliminated the conflict at the RTA waypoint. The conflict was still registered, however, because the LP aircraft were broadcasting their *required* time of arrival at the waypoint, rather than their *estimated* time of arrival. These data illustrate a hazard of broadcasting unachievable flight plan constraints, essentially false intent, in place of true trajectory predictions. Neighboring aircraft may use this false information to make maneuver decisions that, at a minimum, disrupt flight efficiency but may also lead to new conflicts. Broadcasting the commanded trajectory as the intent message, i.e., the four-dimensional path the autoflight system will actually command (assuming no further pilot inputs), would reduce or eliminate the hazards associated with

^{**} Resolutions generated by AOP in the form of modified FMS flight plans that include conflict-free reconnection to the original trajectory

disseminating false information. In fact, a recent update to the ADS-B system performance standards recommends the broadcast of commanded trajectory information^[12].

Who Yielded the Right-of-Way

The objective of using priority flight rules is to reduce the probability of simultaneous maneuvers that may lead to the potentially dangerous dynamic interaction described earlier. It was hypothesized that staggering the alert levels of the two conflicting aircraft would be an effective implementation method for priority flight rules. In this method, the pilot is not directly privy to the actual rule set used or even the resulting relative priority. This approach may have advantages in allowing the use of extensive rule sets that would be too complex for a pilot to remember and for simplifying resolution of multi-conflict situations (i.e., one aircraft in conflict with two or more other aircraft).

The percentages of pilots that maneuvered first in the over-constrained conflict are shown in Figure 5. When priority rules were not in effect, the RH and LH aircraft were equally observed to take the first maneuver action to resolve the conflict, i.e., yield right-of-way. This result indicates that the pilots were not likely applying any previously ingrained right-of-way rules, at least when still distant from the conflicting aircraft. Following this first move, however, considering the constraint conformance data presented earlier, further maneuver interactions between the aircraft occurred in which the RH aircraft frequently prevailed in meeting all constraints. At closer quarters, pilots may therefore be more apt to yield the traditional right-of-way.

When priority rules were in effect, the LP aircraft always yielded right-of-way by maneuvering first. Therefore, the




Figure 5. Pilots that maneuvered first to resolve the over-constrained conflict.

method for implementing the priority system through staggering the conflict alerts successfully induced a bias governing which aircraft will yield.

When Were Constraints Abandoned

The pilots in the simulation were instructed to “notify” the ATS provider as soon as they determined they were not going to meet all of the assigned constraints within the established tolerances. Since the simulation contained no ATS provider position, buttons on the flight deck consoles were provided to represent the capability of sending the message through data link (although voice communication may also be completely acceptable for this action). The pilots were able to declare their inability to meet one or more of the time, position, and/or altitude assignments. The first declaration of “unable” by either pilot in a conflict pair would presumably result in the ATS provider assigning a new set of constraints to that aircraft, and therefore the end of the over-constrained conflict situation as well as any potential hazard therein. The hypothesis was that, with no priority flight rules, the “unable” declaration would occur later than with priority flight rules in effect because it would take longer for the pilots to conclude that meeting the waypoint constraint was incompatible with maintaining traffic separation. The rationale centered on the increased interactions expected between the aircraft as each pilot observed the other aircraft’s maneuver and therefore took additional time to determine whether constraints could still be met.

In considering the mean time when the first pilot in each of the conflict pairs clicked an “unable” button, no noticeable difference was evident when comparing the data runs where priority flight rules were in effect to those where they were not. To gain some insight into whether maneuvering observations may have affected the “unable” declaration time, data is presented in Figure 6 indicating when the first “unable” declaration was made relative to the first maneuver of the opposite aircraft in response to the mutual conflict. When priority flight rules were in effect, the frequency that the first “unable” declaration occurred before the other aircraft maneuvered was observed to increase. Since the priority flight rules caused the HP aircraft to delay maneuvering, more time was simply available to recognize the over-constrained nature of the conflict. Therefore, it is unlikely that the decision to declare “unable” was linked to observed changes in the other aircraft’s state or intent, i.e., interactions between the aircraft were less than originally hypothesized.

Conclusions

An over-constrained traffic conflict in DAG-TM autonomous aircraft operations was studied in a multi-piloted




Figure 6. Comparison of “unable to meet constraints” declaration and the first maneuvers of conflict aircraft.

simulation to determine the need for and utility of priority flight rules for maintaining safety in this potentially hazardous situation. In addition, a proposed scheme for implementing the priority flight rules through staggering the conflict alerting between the two aircraft was tested for effectiveness. The over-constrained conflict scenario investigated involved two piloted aircraft assigned identical en-route waypoint arrival times and altitude crossing restrictions. The study yielded the following conclusions:

- The conflict was found to be safely resolved by the pilots both with and without priority flight rules. Implicit coordination of the resolution though broadcast state and intent data was found to be sufficient for the task.
- Broadcasting intent information consisting of flight plan constraints, rather than true estimates of the actual future trajectory, led to false conflict alerts for the receiving aircraft, triggering unnecessary maneuvering and aircraft interactions close to the waypoint. Broadcasting the commanded trajectory rather than the FMS flight plan trajectory is likely to reduce unnecessary maneuvering and therefore adverse aircraft interactions that may affect predictability and stability in tightly constrained situations.
- The use of priority flight rules had no effect on the percentage of the aircraft population that met all assigned constraints. However, the priority system made completely predictable which aircraft in a given pair would meet the constraints and which aircraft would make the first maneuver to yield right-of-way. Therefore, the proposed scheme for implementing priority flight rules through staggering the conflict alerting between the two aircraft was completely effective. This conclusion has positive implications for the

implementation of a priority flight rule system, should it be determined that one is needed for other reasons. Since the pilot is not required to remember, interpret, or understand the actual relative priority in a given situation, but is only required to follow standard procedures for the given alert level, this scheme permits the application of exceptionally complex rule sets, if needed. Additionally, this scheme may assist in simplifying resolution of multi-conflict situations (i.e., one aircraft in conflict with two or more other aircraft).

The observations from this experiment that such an extraordinary and seemingly hazardous planned conflict situation was safely resolved by observant pilots assisted by appropriate conflict management tools, and that no special restrictions were necessary such as assigning relative priority, strongly support the feasibility of DAG-TM autonomous aircraft operations. This finding suggests that minimal or no involvement by the ATS provider is required to protect safety in conflicts between autonomous aircraft. However, the ATS provider will continue to be a critical component of the traffic flow management for both autonomous and managed aircraft in the DAG-TM concept.

References

- [1] Concept Definition for Distributed Air/Ground Traffic Management (DAG-TM), Version 1.0. NASA Advanced Air Transportation Technologies Project, September 1999.
- [2] Final Report of the RTCA Task Force 3: Free Flight Implementation. RTCA, Inc., Washington D.C., October 1995.
- [3] DAG-TM Concept Element 5 En Route Free Maneuvering Operational Concept Description. Contract NAS2-98005 Research Task Order 72, NASA Advanced Air Transportation Technologies Project, Sept. 2002.
- [4] DAG-TM Concept Element 11 Terminal Arrival: Self-Spacing for Merging and In-trail Separation Operational Concept Description. Contract NAS2-98005 Research Task Order 72, NASA Advanced Air Transportation Technologies Project, December 2002.
- [5] Wing, David J.; Adams, Richard J.; Barmore, Bryan E.; and Moses, Donald: Airborne Use of Traffic Intent Information in a Distributed Air-Ground Traffic Management Concept: Experiment Design and Preliminary Results. Paper 48, 4th USA/Europe ATM R&D Seminar, December 2001
- [6] Wing, David J.; Barmore, Bryan E.; and Krishnamurthy, Karthik: Use of Traffic Intent Information by Autonomous Aircraft in Constrained Operations. AIAA-2002-4555, August 2002.
- [7] Halverson, C. A.: Inside the Cognitive Workplace – New Technology and Air Traffic Control. Ph.D thesis, Univ. of California, San Diego, 1995.
- [8] RTCA SC186: Minimum Operational Performance Standards for 1090 MHz Automatic Dependent Surveillance – Broadcast (ADS-B). RTCA DO-260, September 2000.
- [9] RTCA SC186: Application of Airborne Conflict Management: Detection, Prevention, and Resolution. RTCA DO-263, December 2000.
- [10] Ballin, Mark G.; Sharma, Vivek; Vivona, Robert A.; Johnson, Edward J.; and Ramiscal, Ermin: A Flight Deck Decision Support Tool for Autonomous Airborne Operations. AIAA-2002-4554, August 2002.
- [11] Mondoloni, Stephane; Palmer, Michael T.; and Wing, David J.: Development of a Prototype Airborne Conflict Detection and Resolution Simulation Capability. AIAA-2002-4446, August 2002.
- [12] Barhydt, Richard; and Warren, Anthony W.: Development of Intent Information Changes to Revised Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast (RTCA / DO-242A). NASA/TM-2002-211669, May 2002.

Key Words

DAG-TM, Free Flight, autonomous aircraft, over-constrained conflict, conflict resolution, priority flight rules

Author Biographies

Mr. Wing is a research engineer in the Aviation Operations and Evaluation Branch at the NASA Langley Research Center. Through NASA's Advanced Air Transportation Technologies project, he contributed to the detailed development of the "Distributed Air/Ground Traffic Management" (DAG-TM) concept and research plans, and is currently Langley's Principal Investigator for DAG-TM Concept Element 5 feasibility research. He holds an M.S. in Aeronautical Engineering and a B.S. in Mechanical Engineering, and he is an instrument-rated private pilot.

Dr. Krishnamurthy is a senior ATM research analyst in the Air Traffic Systems Division of The Titan Corporation, currently supporting feasibility research into NASA's DAG-TM concepts and the development of the airborne tools that enable these concepts. He earned a Ph.D. in Aerospace Engineering from Texas A&M University for research related to airborne decision-support tools and artificially intelligent systems for aircraft. He has also dabbled in aircraft design and is a low-time private pilot.

Mr. Barhydt has a B.S. in Aerospace Engineering from the University of Colorado (1995) and an S.M. in Aeronautics and Astronautics from MIT (1997). He is an aerospace engineer at the NASA Langley Research Center, working on the Advanced Air Transportation Technologies project. He has experience in Flight Management System operations and has conducted experiments on the use of intent information on cockpit traffic displays. He is a certified instrument flight instructor and commercial pilot for single engine aircraft. He is also a member of RTCA Special Committee 186, Working Group 4 (Application Technical Requirements) and Working Group 6 (ADS-B MASPS Revision A).

Dr. Barmore is a research engineer in the Aviation Operations and Evaluation Branch at the NASA Langley Research Center. He is currently Langley's Principal Investigator for the "self-spacing for merging and in-trail separation" concept element of NASA's Distributed Air/Ground Traffic Management research project. He holds a Ph.D. in Physics from the College of William and Mary in Virginia and a B.S. from Ohio University.